

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Inventors: Sherif Yacoub, et al. Examiner: Leonard Saint Cyr
Serial No.: 10/668,141 Group Art Unit: 2626
Filed: September 23, 2003 Docket No.: 200300101-I
Title: System and Method Using Multiple Automated Speech Recognition
Engines

APPEAL BRIEF UNDER 37 C.F.R. § 41.37

Mail Stop Appeal Brief - Patents
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Sir:

This Appeal Brief is filed in response to the Final Office Action mailed December 8, 2008 and Notice of Appeal filed on March 9, 2009.

AUTHORIZATION TO DEBIT ACCOUNT

It is believed that no extensions of time or fees are required, beyond those that may otherwise be provided for in documents accompanying this paper. However, in the event that additional extensions of time are necessary to allow consideration of this paper, such extensions are hereby petitioned under 37 C.F.R. § 1.136(a), and any fees required (including fees for net addition of claims) are hereby authorized to be charged to Hewlett-Packard Development Company's deposit account no. 08-2025.

I. REAL PARTY IN INTEREST

The real party in interest is Hewlett-Packard Development Company, LP, a limited partnership established under the laws of the State of Texas and having a principal place of business at 20555 S.H. 249 Houston, TX 77070, U.S.A. (hereinafter "HPDC"). HPDC is a Texas limited partnership and is a wholly-owned affiliate of Hewlett-Packard Company, a Delaware Corporation, headquartered in Palo Alto, CA. The general or managing partner of HPDC is HPQ Holdings, LLC.

II. RELATED APPEALS AND INTERFERENCES

There are no known related appeals, judicial proceedings, or interferences known to appellant, the appellant's legal representative, or assignee that will directly affect or be directly affected by or have a bearing on the Appeal Board's decision in the pending appeal.

III. STATUS OF CLAIMS

Claims 1 – 20 are pending in the application and stand finally rejected. The rejection of claims 1 – 20 is appealed.

IV. STATUS OF AMENDMENTS

No amendments were made after receipt of the Final Office Action. All amendments have been entered. Further, no amendments were made to the claims during prosecution. As such, all claims are original.

V. SUMMARY OF CLAIMED SUBJECT MATTER

The following provides a concise explanation of the subject matter defined in each of the claims involved in the appeal, referring to the specification by page and line number and to the drawings by reference characters, as required by 37 C.F.R.

§ 41.37(c)(1)(v). Each element of the claims is identified by a corresponding reference to the specification and drawings where applicable. Note that the citation to passages in the specification and drawings for each claim element does not imply that the limitations from the specification and drawings should be read into the corresponding claim element or that these are the sole sources in the specification supporting the claim features.

Claim 1

A method of automatic speech recognition (ASR), comprising (Figure 3 shows a method of automated speech recognition using plural different ASR engine. The method is described in connection with Figure 1: see paragraph [0031] on p. 8):

receiving a speech utterance from a user (Fig. 3, #300: A participant or user (such as a telephone caller) telephones or otherwise establishes communication between communication device 40 and communication network 10. Per block 300, the communication device provides communication network 10 with an electronic input signal in a digital format. See paragraph [0031] on p. 8.);

assessing resources of each of a plurality of different ASR engines to determine which of the plurality of different ASR engines are busy serving users (Fig. 3, #310: Per block 310, the host computer 50 assesses the resources of the system. At this stage, for example, ports of different ASR engines 60, memory 90, database 100, or processing power of CPU 80 can be evaluated for availability. See paragraph [0032] on p. 8. Per blocks 320 and 330, a determination is made which ASR engines are busy: see paragraph [0033] and [0034] on p. 8.);

assigning the speech utterance to a single ASR engine when the plurality of different ASR engines are busy such that assessing resources is within a threshold value (Fig. 3, #330: Per block 330, if the system is busy, the host computer 50, in cooperation with the resource management application 110, assigns the input signal to a single ASR engine. See paragraph [0034] on p. 8.);

assigning the speech utterance to a plurality of different ASR engines when the plurality of different ASR engines are not busy such that assessing resources is within a threshold value (Fig. 3, # 350: On the other hand, per block 350, if the system is not busy, the host computer, in cooperation with the resource management application 110, assigns the input signal to multiple ASR engines. Here, the recognized text from the selected ASR engines is combined to yield more accurate recognized text when compared to a single ASR engine. See paragraph [0035] on p. 8.); and

generating text of the speech utterance with either the single ASR engine or plurality of ASR engines (Fig. 3, #340 and 360: The assigned ASR engine or engines generate the recognized text of the input signal. See paragraphs [0034] and [0035] on p. 8.).

Claim 3

The method of claim 1 wherein assessing resources further comprises evaluating processing power (The host computer assesses resources of the system, and these resources include processing power of the CPU: see paragraphs [0032] on p. 8 and [0045] on p. 11.).

Claim 4

The method of claim 1 wherein assessing resources further comprises monitoring memory utilization and input/output utilization (The host computer assesses resources of the system, and these resources include free memory and I/O usage: see paragraphs [0032] on p. 8 and [0045] on p. 11.).

Claim 8

An automatic speech recognition (ASR) system comprising (FIG. 1 shows three ASR systems 60A, 60B, and 60C. FIG. 2 shows a block diagram of an exemplary embodiment of an ASR system 60A: see paragraph [0021].):

means for processing a digital input signal from an utterance of a user (Example means is CPU 80 in Fig. 1. Per block 300 in Fig. 3, a participant or user (such as a telephone caller) telephones or otherwise establishes communication between

communication device 40 and communication network 10. Per block 300, the communication device provides communication network 10 with an electronic input signal in a digital format. See paragraph [0031].);

means for evaluating resources of the ASR system to determine whether the ASR system is busy processing utterances of users (Example means is host computer 50 in Fig. 1. The host computer 50 assesses the resources of the system. Per block 310 of Fig. 3, the host computer 50 assesses the resources of the system. At this stage, for example, ports of different ASR engines 60, memory 90, database 100, or processing power of CPU 80 can be evaluated for availability. See paragraph [0032] on p. 8.); and

means for selecting between a single ASR engine and a group of ASR engines to recognize the utterance of the user, wherein the means for selecting utilizes the evaluation of resources to select the single ASR engine when the ASR system is busy processing the utterances of the users and to select the group of ASR engines when the ASR system is not busy processing the utterances of the users (Example means is host computer 50 in Fig. 1. The host computer 50 selects between a single ASR engine or multiple different ASR engines. Per block 330 in Fig. 3, if the system is busy, the host computer 50, in cooperation with the resource management application 110, assigns the input signal to a single ASR engine. Per block 350 in Fig. 3, if the system is not busy, the host computer, in cooperation with the resource management application 110, assigns the input signal to multiple ASR engines. See paragraphs [0034] and [0035] on p. 8.).

Claim 10

The ASR system of claim 9 wherein the means for evaluating resources of the system also monitors available processing power of the system (The host computer assesses resources of the system, and these resources include processing power of the CPU: see paragraphs [0032] on p. 8 and [0045] on p. 11.).

Claim 14

A system, comprising (Fig. 1 shows a diagram of a voice telephone system, and Fig. 2 shows a diagram of an ASR system: see paragraphs [0020] and [0021] on p. 5.):

a computer system comprising a central processing unit coupled to a memory and resource management application (Fig. 1 shows a host computer system 50 comprising a CPU 80, memory 90, extracted algorithm 110, and resource management application 110 coupled through buses 120. See paragraph [0019]); and

a plurality of different automatic speech recognition (ASR) engines coupled to the computer system, wherein the computer system assesses resources being used by each of the plurality of different ASR engines and selects a single ASR engine to analyze a speech utterance when the system is busy and selects multiple ASR engines to analyze the speech utterance when the system is not busy (Fig. 1 shows plural ASR systems 60A, 60B, and 60C coupled to the computer system 50. See paragraph [0018] on p. 4. The host computer 50 selects between a single ASR engine or multiple different ASR engines. Per block 330 in Fig. 3, if the system is busy, the host computer 50, in cooperation with the resource management application 110, assigns the input signal to a single ASR engine. Per block 350 in Fig. 3, if the system is not busy, the host computer, in cooperation with the resource management application 110, assigns the input signal to multiple ASR engines. See paragraphs [0034] and [0035] on p. 8.).

Claim 15

The system of claim 14 wherein the computer system selects an ASR engine that has most available resources (Various resources are examined to determine an engine with the most available resources: see paragraphs [0032] on p. 8, [0043] on p. 10, and [0045] on p. 11.).

VI. GROUNDS OF REJECTION TO BE REVIEWED ON APPEAL

Claims 1 – 20 are rejected under 35 USC § 103(a) as being unpatentable over US publication number 2002/0193991 (Bennett) in view of USPN 6,728,671 (Johnson) and USPN 4,641,342 (Watanabe).

VII. ARGUMENT

The rejection of claims 1 – 20 is improper, and Appellants respectfully request reversal of these rejections.

The claims do not stand or fall together. Instead, Appellants present separate arguments for various claims. Each of these arguments is separately argued below and presented with separate headings and sub-heading as required by 37 C.F.R. § 41.37(c)(1)(vii).

Claim Rejections: 35 USC § 103(a)

Claims 1 – 20 are rejected under 35 USC § 103(a) as being unpatentable over US publication number 2002/0193991 (Bennett) in view of USPN 6,728,671 (Johnson) and USPN 4,641,342 (Watanabe). These rejections are traversed.

Principles of Law: Claim Construction

During examination of a patent application, pending claims are given their broadest reasonable construction consistent with the specification (see *In re Prater*, 415 F.2d 1393, 1404-05 (CCPA 1969); *In re Am. A cad. a/Sci.Tech Ctr.*, 367 F.3d 1359, 1364 (Fed. Cir. 2004)).

Although a patent applicant is entitled to be his or her own lexicographer of terms in a claim, in *ex parte* prosecution the lexicography must be within limits. *In re Carr*, 347 F.2d 578,580 (CCPA 1965). The applicant must do so by placing such definitions in the specification with sufficient clarity to provide a person of ordinary skill in the art with clear and precise notice of the meaning that is to be construed. See also *In re Paulsen*, 30 F.3d 1475, 1480 (Fed. Cir. 1994) (although an inventor is free to define the specific terms used to describe the invention, this must be done with reasonable clarity, deliberateness, and precision; where an inventor chooses to give terms uncommon meanings, the inventor must set out any uncommon definition in some manner within the patent disclosure so as to give one of ordinary skill in the art notice of the change).

Principles of Law: Obviousness

The test for determining if a claim is rendered obvious by one or more references for purposes of a rejection under 35 U.S.C. § 103 is set forth in *KSR International Co. v. Teleflex Inc.*, 550 U.S._, 82 USPQ2d 1385 (2007):

Under §103, the scope and content of the prior art are to be determined; differences between the prior art and the claims at issue are to be ascertained; and the level of ordinary skill in the pertinent art resolved. Against this background the obviousness or nonobviousness of the subject matter is determined. Such secondary considerations as commercial success, long felt but unsolved needs, failure of others, etc., might be utilized to give light to the circumstances surrounding the origin of the subject matter sought to be patented. Quoting *Graham v. John Deere Co. of Kansas City*, 383 U.S. 1 (1966).

As set forth in MPEP 2143.03, to ascertain the differences between the prior art and the claims at issue, “[a]ll claim limitations must be considered” because “all words in a claim must be considered in judging the patentability of that claim against the prior art.” *In re Wilson*, 424 F.2d 1382, 1385.

According to the Examination Guidelines for Determining Obviousness Under 35 U.S.C. 103 in view of *KSR International Co. v. Teleflex Inc.*, Federal Register, Vol. 72, No. 195, 57526, 57529 (October 10, 2007), once the *Graham* factual inquiries are resolved, there must be a determination of whether the claimed invention would have been obvious to one of ordinary skill in the art based on any one of the following proper rationales:

- (A) Combining prior art elements according to known methods to yield predictable results;
- (B) Simple substitution of one known element for another to obtain predictable results;
- (C) Use of known technique to improve similar devices (methods, or products) in the same way;
- (D) Applying a known technique to a known device (method, or product) ready for improvement to yield predictable

results; (E) “Obvious to try”—choosing from a finite number of identified, predictable solutions, with a reasonable expectation of success; (F) Known work in one field of endeavor may prompt variations of it for use in either the same field or a different one based on design incentives or other market forces if the variations would have been predictable to one of ordinary skill in the art; (G) Some teaching, suggestion, or motivation in the prior art that would have led one of ordinary skill to modify the prior art reference or to combine prior art reference teachings to arrive at the claimed invention. *KSR International Co. v. Teleflex Inc.*, 550 U.S._, 82 USPQ2d 1385 (2007).

Furthermore, as set forth in *KSR International Co. v. Teleflex Inc.*, quoting from *In re Kahn*, 441 F.3d 977, 988 (CA Fed. 2006), “[R]ejections on obviousness grounds cannot be sustained by mere conclusory statements; instead, there must be some articulated reasonings with some rational underpinning to support the legal conclusion of obviousness.”

Therefore, if the above-identified criteria and rationales are not met, then the cited reference(s) fails to render obvious the claimed invention and, thus, the claimed invention is distinguishable over the cited reference(s).

Scope and Content of Art and Overview of Claims

As a precursor to the arguments, Appellants provide an overview of the claims and the primary references (Bennett, Johnson, and Watanabe). This overview will assist in determining the scope and content of the prior art as required in Graham (see *Graham v. John Deere Co. of Kansas City*, 383 U.S. 1, 17-18 setting out an objective analysis for applying 103 rejections).

By way of background, automated speech recognition (ASR) engines recognize speech and then perform tasks without the use of additional human intervention. ASR systems and networks that utilize a single ASR engine are challenged to recognize accurately and consistently various speech patterns and utterances.

The claims are directed to ASR methods and systems. When a speech utterance is received, the available resources of plural different ASR engines are assessed to

determine which ASR engines are busy serving users. Based on the assessed resources, the speech utterance is assigned to a single ASR engine when the ASR engines are busy. Alternatively, the speech utterance is assigned to multiple different ASR engines when the multiple engines are not busy.

Bennett teaches an ASR system in which the input stream is routed to one or more speech recognition engines. In some situations, Bennett does not route the input stream to all of the speech recognition engines. This situation occurs when a speech recognizer is identified as being only for command and control.

Johnson teaches an ASR system that increases ASR capacity without increasing a number of current ASR input channels. Johnson uses delay functionalities to provide this increase in ASR capacity. Thus, Johnson assesses usage of the ASR input channels to determine a level of delay.

Watanabe teaches a voice input system that avoids a misrecognition or a rejection when a user speaks before the input cue is given. In this system, the number of the voice recognition processing means is reduced below the number of the user channels to simplify the hardware structure.

Differences Between the Art and Claims

Claims 1 – 20 recite one or more elements that are not taught or suggested in Bennett in view of Johnson and Watanabe. These missing elements show that the differences between the combined teachings in the art and the recitations in the claims are great. As such, the pending claims are not a predictable variation of the art to one of ordinary skill in the art.

These differences are shown below and presented with separate headings for different claim groups.

Sub-Heading: Claims 1 – 2, 5 – 9, 11 – 14, and 16 – 20

Independent claim 1 is selected for discussion.

As one example, independent claim 1 recites assessing resources of each of a plurality of different ASR engines to determine which of the plurality of different ASR engines are busy serving users. This claim element is not taught in Bennett in view of Johnson and Watanabe.

In Bennett, the input stream is routed to one or more speech recognition engines (see Bennett at paragraph [0023]). In some situations, Bennett does not route the input stream to all of the speech recognition engines. Bennett explains when this situation occurs:

For example, the incoming stream may be at a point in the dialogue with the system that is beyond command and control. The speech recognition system may therefore not send the incoming stream to the recognizers that have been identified by the system as being only for command and control. Other types of cues may also be available to the system allowing it to route the stream to a subset of recognizers. (See Bennett at paragraph [0023]).

Thus, Bennett teaches that the incoming stream is not routed to speech recognition engines that are identified as being only for command and control. Bennett, however, does not assess resources of each of a plurality of different ASR engines to determine which of the plurality of different ASR engines are busy serving users. Instead, Bennett identifies which ASR engine is for command and control. Furthermore, although Bennett does state that other types of cues may be available, Bennett never suggests that such cues include assessing resources of each of a plurality of different ASR engines to determine which of the plurality of different ASR engines are busy serving users. Bennett also does not teach or suggest assigning the speech utterance to a single ASR engine when the plurality of different ASR engines are busy and assigning the speech utterance to a plurality of different ASR engines when the plurality of different ASR engines are not busy.

Johnson teaches an ASR system that increases ASR capacity without increasing a number of current ASR input channels (see Johnson column 3, lines 25-29). Johnson uses delay functionalities to provide this increase in ASR capacity: “[B]ased upon the degree of utilization of the ASR 240 input channels, the processor 210 may select various delay modes, such as no delay, light (low) delay, intermediate (medium) delay, or heavy (high) delay” (see Johnson at column 6, lines 26-29). The use of a delay mode is reiterated in the method of Fig. 2 in Johnson (see column 8, lines 33-50). Here, Johnson explains that the delay mode uses “any various combinations of increased periods of silence and longer durations of messages” (column 8, lines 46-47).

Thus, Johnson assesses usage of the ASR input channels and then uses delays. By contrast, claim 1 recites assessing resources of plural different ASR engines to determine which of the engines are busy serving users. The claim then recites assigning the speech utterance to a single ASR engine when the plurality of different ASR engines are busy and assigning the speech utterance to a plurality of different ASR engines when the plurality of different ASR engines are not busy. Johnson assesses ASR input channels and uses delays. Claim 1 assesses ASR engine resources and assigns the speech utterance to a single ASR engine when the ASR engines are busy and to plural ASR engines when the ASR engines are not busy. Johnson does not assess resources of plural different ASR engines to perform such an assignment to a single or multiple ASR engines. Johnson assesses resources to determine a level of delay.

Clearly, Bennett in view of Johnson does not teach or even suggest the following claim elements:

assigning the speech utterance to a single ASR engine when the plurality of different ASR engines are busy such that assessing resources is within a threshold value;

assigning the speech utterance to a plurality of different ASR engines when the plurality of different ASR engines are not busy such that assessing resources is within a threshold value.

As discussed below, the Examiner agrees with this conclusion.

Admission of Examiner and Response to Watanabe

In the Final Office Action, the Examiner admits the following:

However, Bennett et al., in view of Johnson et al., do not specifically teach assigning the speech utterance to a single ASR engine when the plurality of different ASR engines are busy such that assessing resources is within a threshold value; assigning the speech utterance to a plurality of different ASR engines when the plurality of different ASR engines are not busy such that assessing resources is within a threshold value. {see Final OA at p. 4}.

Appellants agree with this admission. The Examiner, however, attempts to cure this deficiency with the Watanabe and cites the following location:

Watanabe et al., teach detector can also be used to selectively switch an active one of a plurality of user channels to one of a smaller number of voice recognizers. The control means searches for any idle input terminal 121- 12n at the recognition means ... when an idle terminal is found, the change-over switch 14 operates to connect the requested user channel with the idle recognition input terminal. A change-over switch 14 connects m voice pattern outputs to n separate recognition input terminals 121, 122, . . . , 12j, . . . , 12n of a recognition means 15 which recognizes the voices received on its n separate recognition input terminals as the outputs from the change-over switch 14. The recognition means 15 can recognize n separate voice simultaneously (Abstract, lines 7 – 10; col.4, lines 5 – 12, and 36 – 46; col.6, lines 4 – 6). {See Final OA at p. 4}.

Appellants respectfully submit that the Examiner is mischaracterizing the teachings in Watanabe and taking such teachings out of context.

Watanabe is solving a problem that occurs when users speak before the system is ready to receive their voice (for example, a user speaks before an input cue, such as a

beep, is provided to the user: see column 1, lines 11 – 22). Watanabe solves this problem by re-issuing an input cue signal (for example, a beep) when the user speaks before the system is ready (see column 2, lines 9 – 13). Further, Watanabe uses a change over switch to ensure that only the user channel requesting the voice input is connected to the recognition. Watanabe explains how this system works with multiple users simultaneously using the system:

Namely, each user utilizes a voice input and voice output, so that, while the voice output is being issued, the recognition means is not used. This period is idle time for the recognition unit. In order to make use of this idle time for the processing of voice input from other user channel, a change-over switch is provided between the user's circuit and the recognition input terminals so that only the user channel requesting the voice input is connected to the recognition means {see column 3, lines 28 – 36}.

With this understanding, Appellants respectfully assert that the Examiner is mischaracterizing the teachings in Watanabe. The section of Watanabe quoted by the Examiner teaches that the detector searches for an idle input terminal and then switches to this idle input. The change-over switch is used to ensure only the user channel requesting the voice input is connected to the recognizer.

The quoted sections of Watanabe do not teach or even suggest assigning a speech utterance to a single ASR engine when a plural different ASR engines are busy such that assessing resources is within a threshold value. The quoted sections of Watanabe also do not teach or even suggest assigning the speech utterance to a plurality of different ASR engines when the plurality of different ASR engines are not busy such that assessing resources is within a threshold value. As explained above, Watanabe is using a very different system for a very different purpose.

The differences between the claims and the teachings in the art are great since the references fail to teach or suggest all of the claim elements. As such, the pending claims are not a predictable variation of Bennett in view of Johnson and Watanabe to one of ordinary skill in the art.

For at least these reasons, the claims are allowable over the art of record.

Sub-Heading: Dependent Claims 3 and 10

Claims 3 and 10 recite that assessing resources comprises evaluating processing power to determine which ASR engines are busy. The Examiner argues that this claim element is taught in Johnson at column 5, lines 26-28. Appellants respectfully disagree.

Column 5, lines 26-28 in Johnson teaches that an observer module detects power on input channels “indicating that a caller has begun to speak and, therefore, that the corresponding input channel is to be processed for speech recognition by the ASR 240.” By contrast, claim 1 recites that the assessing occurs to determine which of plural different ASR engines are busy serving users. Johnson, however, examines power to determine if the caller started to speak. Power is being used in two very different functions in claims 3 and 10 and Johnson.

Sub-Heading: Dependent Claim 4

Claim 4 recites that assessing the resources comprises monitoring memory utilization and input/output utilization to determine which ASR engines are busy. The Examiner argues that this claim element is taught in Johnson at column 2, lines 23-25. Appellants respectfully disagree.

Column 2, lines 23-25 in Johnson teaches that input is stored in a memory buffer. This teaching is not related to using memory to assess resources. Claim 1 recites that the assessing occurs to determine which of plural different ASR engines are busy serving users. Johnson, however, is using memory merely to store input, not to assess resources to determine whether an ASR engine is busy.

Sub-Heading: Dependent Claim 15

Claim 15 recites that the computer system selects an ASR engine that has the most available resources. The Examiner argues that Bennett teaches this element at paragraph [0020], lines 7 – 9. Appellants respectfully disagree.

Paragraph [0020] in Bennett discusses the function of the recognizers. The recognizers convert speech to text and have different performance levels. This difference

in performance depends on which manufacturer made the recognizer: A recognizer from one manufacturer can have a higher accuracy than a recognizer from a different manufacturer.

Bennett does not select an ASR engine or recognizer with the most available resources as recited in claim 15. Bennett does not determine which recognizer has available resources. Instead, Bennett merely acknowledges that some recognizers have better accuracy than others. Knowing an accuracy of a recognizer is very different than selecting a recognizer that has the most available resources.

CONCLUSION

In view of the above, Appellants respectfully request the Board of Appeals to reverse the Examiner's rejection of all pending claims.

Any inquiry regarding this Amendment and Response should be directed to Philip S. Lyren at Telephone No. 832-236-5529. In addition, all correspondence should continue to be directed to the following address:

Hewlett-Packard Company
Intellectual Property Administration
P.O. Box 272400
Fort Collins, Colorado 80527-2400

Respectfully submitted,

/Philip S. Lyren #40,709/

Philip S. Lyren
Reg. No. 40,709
Ph: 832-236-5529

VIII. Claims Appendix

1. A method of automatic speech recognition (ASR), comprising:
 - receiving a speech utterance from a user;
 - assessing resources of each of a plurality of different ASR engines to determine which of the plurality of different ASR engines are busy serving users;
 - assigning the speech utterance to a single ASR engine when the plurality of different ASR engines are busy such that assessing resources is within a threshold value;
 - assigning the speech utterance to a plurality of different ASR engines when the plurality of different ASR engines are not busy such that assessing resources is within a threshold value; and
 - generating text of the speech utterance with either the single ASR engine or plurality of ASR engines.
2. The method of claim 1 wherein assessing resources further comprises monitoring port utilization for each ASR engine.
3. The method of claim 1 wherein assessing resources further comprises evaluating processing power.
4. The method of claim 1 wherein assessing resources further comprises monitoring memory utilization and input/output utilization.

5. The method of claim 1 wherein assessing resources further comprises monitoring a number of users providing speech utterances.
6. The method of claim 1 wherein assigning the speech utterance to a single ASR engine if assessing resources is within a threshold value occurs when port utilization of the single ASR engine is lower than a port utilization threshold of about 80%.
7. The method of claim 1 wherein assigning the speech utterance to a plurality of different ASR engines if assessing resources is within a threshold value occurs when port utilization of two ASR engines is lower than a predefined threshold of about 75%.
8. An automatic speech recognition (ASR) system comprising:
 - means for processing a digital input signal from an utterance of a user;
 - means for evaluating resources of the ASR system to determine whether the ASR system is busy processing utterances of users; and
 - means for selecting between a single ASR engine and a group of ASR engines to recognize the utterance of the user, wherein the means for selecting utilizes the evaluation of resources to select the single ASR engine when the ASR system is busy processing the utterances of the users and to select the group of ASR engines when the ASR system is not busy processing the utterances of the users.
9. The ASR system of claim 8 wherein the means for evaluating resources of the system monitors port utilization of the ASR engines to determine when an ASR engine is busy.

10. The ASR system of claim 9 wherein the means for evaluating resources of the system also monitors available processing power of the system.

11. The ASR system of claim 8 further comprising a means for combining results of ASR engines if the group of ASR engines is selected, the group of ASR engines being adapted to provide a more accurate recognition of the utterance than a single ASR engine.

12. The ASR system of claim 8 wherein the means for evaluating resources of the system evaluates resources to simultaneously run multiple ASR engines.

13. The ASR system of claim 8 wherein the means for evaluating resources of the system evaluates ASR ports, system resources, and call handlers.

14. A system, comprising:

a computer system comprising a central processing unit coupled to a memory and resource management application; and

a plurality of different automatic speech recognition (ASR) engines coupled to the computer system, wherein the computer system assesses resources being used by each of the plurality of different ASR engines and selects a single ASR engine to analyze a speech utterance when the system is busy and selects multiple ASR engines to analyze the speech utterance when the system is not busy.

15. The system of claim 14 wherein the computer system selects an ASR engine that has most available resources.
16. The system of claim 14 further comprising a telephone network comprising at least one switching service point coupled to the computer system.
17. The system of claim 16 further comprising at least one communication device in communication with the switching service point to provide the speech utterance.
18. The system of claim 14 wherein the resource management application comprises a recognition proxy component and a resource monitoring component.
19. The system of claim 18 wherein the resource management component collects and analyzes information about the resources available on the system.
20. The system of claim 19 wherein the resource monitoring component mediates between the plurality of ASR engines and the resource management component.

IX. EVIDENCE APPENDIX

None.

X. RELATED PROCEEDINGS APPENDIX

None.